ANL-24/20

ECP Workflows Survey

Data Science and Learning

ANL-24/20

ECP Workflows Survey Report

prepared by
Justin M Wozniak
Data Science & Learning Division, Argonne National Laboratory

April 30, 2017

ECP Workflow Survey Report 2017

Justin M. Wozniak <woz@anl.gov>

April 26, 2017

Contents

1 Overview

The ECP applications teams were invited to participate in an online survey from April 10-17, 2017. The survey was hosted on Google Forms and consisted of questions determined by representatives from ECP Software Technology teams and the CODAR codesign center.

The introductory text for the survey was as follows:

> Thank you for participating in the ECP Workflows Survey for 2017. This survey will help ECP leadership determine directions in workflows development and deployment on exascale systems.
>
> For the purposes of this survey, a workflow is a collection of computational tasks and operations in support of some greater scientific goal. This involves the usage of programs and libraries, including their installation, access, execution, and monitoring. It also involves the management of data, its creation, access, sharing, movement, and life cycles. Workflows involve human and automated interactions both on and off of large-scale computers, as well as among scientific codes and system services, on the login nodes, compute nodes, or combinations. Information about the characteristics of these items will help in the development of workflow systems and the overall exascale ecosystem.
>
> Most questions in the survey have multiply-selectable checkboxes so you can easily provide more than one response, fill in a new answer, or indicate that you don't know. You can also use the follow-up comment fields for further information that you may wish to add, which will help us interpret the survey results. For example, if you have multiple workflows that may be of interest, you may use multiple checkboxes and describe your answers, or take the survey twice, indicating the different workflow of interest in your project description.

Unless otherwise noted, all questions were multiple choice, and multiple selections could be made (including the numerical choices). This is because the various aspects of a workflow or various workflows within a project could lead to different answers. Participants were also able to specify "Other" and write their own response. Each question was followed up with a corresponding free-form comment text field.

Responses are reported here in three formats:

- A heavily processed and annotated version (§2)
- A readable but expansive full-text representation (§3)
- A Google-generated report (§4)

The raw data is also available in its original form, a Google Sheet, and as an exported CSV file. Comment text is only visible in the full-text and Google reports.

The survey itself is appended to the end of this document (§5).

The survey received 13 responses from the 22 application development projects. The organizers would like to thank each of them for participating. The participants are as follows, in order of survey entry time:

Name	Area	Project
Mark Taylor	Climate	ACME-MMF
Rajeev Jain	Urban science	Urban
Amedeo Perazzo	X-ray science	ExaFEL
Mark Gordon	Chemistry	GAMESS
Anders Petersson	Earthquakes	Seismic risk simulations
Balint Joo	Physics	Lattice QCD
Jean-Luc Vay	Physics	WarpX
Jaron Krogel	Physics	QMCPACK
Ray Bair	Chemistry	NWChemEx
Jibo Sanyal	Urban science	Urban
Salman Habib	Cosmology	ExaSky
Kathy Yelick	Biology	ExaBiome
Mitchell Wood	Materials science	EXAALT

Note that two entries were made for the Urban science project by different persons.

2 Responses: Tabulated

2.1 How many programs or software components may be brought together in one of your workflows?

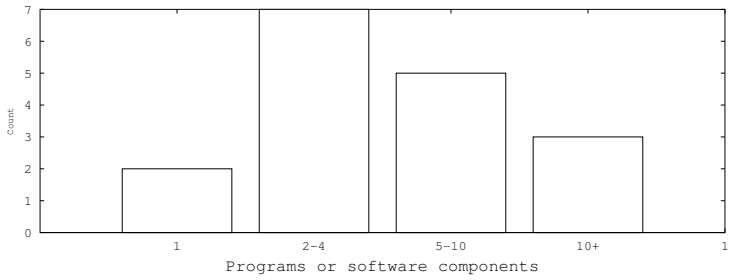

Interpretation: Roughly half of the respondent workflows combine 5 or more programs or software components, which is a significant software management challenge.

2.2 What size jobs make up your workflows?

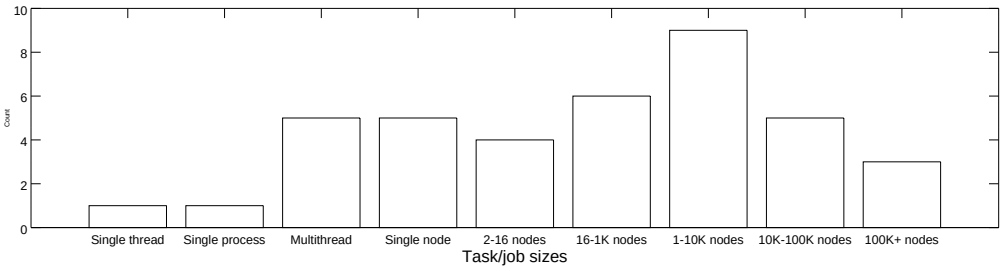

Interpretation: Roughly half of the respondent task sizes indicate that multi-node, yet not full-system, tasks make up their workflows.

2.3 How long are the tasks in your workflow?

Count	Response
8	12+ hours
6	16 minutes - 1 hour
6	1 hour - 12 hours
5	100 seconds - 15 minutes
1	from seconds to tens of minutes
1	Highly variable
1	1-99 seconds

Interpretation: The respondents are primarily interested in moderately long to very long tasks, although there is a presence of tasks that run for less than 15 minutes, which would pose challenges to traditional system schedulers.

2.4 In a given process, how many threads might you use efficiently?

Count	Response
5	64+
4	2-7
2	8-63
2	1
1	as many as the cores support
1	Nek5000 threaded version is under development.
1	I don't know
1	2-64

Interpretation: The respondents indicate good readiness for many-core (64+ core) architectures.

2.5 How much memory can you use effectively on a single node?

Count	Response
7	64+ GB
4	16-63 GB
3	8-15 GB
3	2-7 GB
2	I don't know
1	1 GB or less

Interpretation: The respondents indicate that large node memories (64+ GB) are required.

2.6 What accelerators are your code capable of using?

Count	Response
11	Xeon Phi
10	GPU
1	FPGA
1	ARM

Interpretation: The respondents indicate roughly equal readiness for many-core and GPU architectures; no respondents selected the provided option "None."

2.7 How many tasks (program or software component invocations) can be part of a single workflow?

Count	Response
5	2-19
5	1
4	20-99
3	10,000-999,999
2	100-999
1	This depends on task definition; the question isn't clear. In some sense the number of independent tasks in the code is set by the number of MPI ranks; this follows from how the code is designed. However for load balancing purposes we use independent work slabs which are much smaller than the 1-rank workload.
1	Currently we run 1 at a time but in production assembly we would want a workflow to push through many jobs. (There are millions to process.)
1	1000-9999

Interpretation: The respondents are primarily interested in workflows that contain <100 tasks, but there is some interest in workflows with many (thousands) tasks.

2.8 How much data is consumed and produced by individual tasks in your workflow?

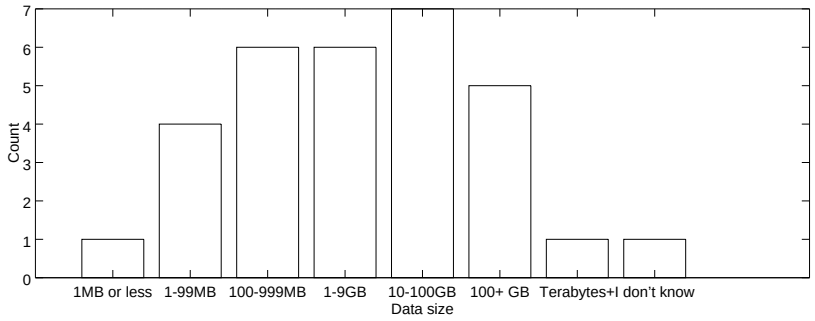

2.9 What use cases are you considering for burst buffers?

Count	Response
6	Implicitly - behind some I/O library
5	Implicitly - the workflow system should be able to stage my data
4	Implicitly - behind the POSIX interfaces
4	Explicitly - do my own buffering/staging
1	We have not started playing with burst buffers
1	My workflow cannot benefit from burst buffers.
1	I don't know what burst buffers are.
1	CEED and ExaHDF5 dependency

Interpretation: A large number of respondents indicate that the workflow system, or a hand-crafted workflow, should make direct use of burst buffers.

2.10 What types of code coupling are required by your workflow?

Count	Response
9	Communication via files
7	Explicit MPI communication
6	zero-copy
6	Shared-memory
3	Databases
3	Data exchange via other communication library or framework
1	my tasks are independent.
1	None

Interpretation: Many reported use cases still plan to use files for intermediate workflow products, although there is some interest in advanced technologies.

2.11 What workflow systems do you currently use?

Rationale: This was a free-form text field to gauge interest in existing workflow systems.

Count	Response
1	iRODS
1	custom
1	Swift and scripts
1	Still collection data/setting-up individual physics simulations
1	Scripts (PERL/Python/Shell)
1	None in the main code; galaxy is used for an offline data analytics service
1	Nexus (custom Python packaged w/ qmcpack)
1	HipMer (the non-metagenome version) has been integrated into KBase
1	Batch script; sometimes written from python script.
1	Ad hoc

2.12 What workflow systems are you considering using?

Rationale: This was a free-form text field to gauge interest in existing workflow systems.

Responses: 1 respondent indicated Swift, the others were unsure of available solutions.

2.13 What paradigms drive your workflows?

Count	Response
7	Simulation + analytics
6	Coupled simulations (e.g. multi-physics multi-scale)
6	Bag of tasks
3	Machine learning
2	Optimization
1	We are still in early stages of simulation setup
1	Uncertainty quantification

Interpretation: Simulation-simulation and simulation-analytics coupling are the most prevalent workflow use cases, although there is some interest in more advanced patterns.

2.14 Do your workflows have a human-in-the-loop?

Rationale: This question intended to gauge community interest in fully-automated workflows versus workflows with human intervention.

Count	Response
8	Somewhere in between
5	The workflow is fully automated
1	I don't know
1	Humans have to inspect every task

2.15 What external data impacts your workflow?

Count	Response
5	Streams from other computational jobs
4	Experiment (e.g. light sources telescopes)
2	I don't know

Interpretation: There is significant interest in streaming data sources.

2.16 What is your preferred syntax for workflow development?

Count	Response
10	Embedded in a scripting language (Python etc.)
6	Embedded in a compiled language (C/C++/Fortran)
3	Need a standard programming language (shell/Makefile)
2	Static data-flow workflow description (Directed Acyclic Graph of Processes/Tasks)
2	I don't know
1	Visual workflow composition (graphical user interface GUI)
1	Jupyter (on which KBase is built)
1	Bag of tasks

2.17 Does your workflow have interoperability requirements?

Count	Response
7	Need to integrate with MPI
4	Need to integrate with the system scheduler
1	scripts
1	Need to integrate with GASNet/UPC
1	?

2.18 How are your workflows composed?

Count	Response
12	My codes must/may be separate executables
5	My codes must/may be executed as libraries in a framework
2	My codes must/may be run on the same node

2.19 What are your preferred interfaces for managing simulations and human-in-the-loop interactions?

Count	Response
10	Standard command line interface
4	Notebooks: interactive sharable scripting
3	Other GUI dashboard or web interface
2	I don't know

2.20 How important is the reproducibility of workflows?

Rationale: This question asked for an importance measurement on a scale of 1-10, with 10 most important.

Count	Response
7	10
5	8
1	9

2.21 How important is robust restart on failure?

Rationale: This question asked for an importance measurement on a scale of 1-10, with 10 most important.

Count	Response
5	10
4	8
1	9
1	7
1	5
1	1

2.22 Can your entire workflow be checkpointed/restarted?

Count	Response
6	Yes
4	In part
3	No
1	I don't know

2.23 Which of the following describes your provenance needs or capabilities?

Count	Response
6	I collect my own provenance data
3	Log files are fine
3	I would like to use a database system
3	I also need integration with performance analysis tools
2	I don't know what provenance is
1	KBase
1	I need workflow provenance data

2.24 How many files are accessed by the whole workflow?

Count	Response
9	11-10,000 in total
3	1-10 in total
2	11-100 per task
2	1-10 per task
1	We'll stream the ingress data
1	I don't know

3 Responses: Full text

The raw survey data is appended here.

4 Responses: Google report

The Google-provided report is appended here.

5 Survey questions

The Google-formatted survey is appended here.

ECP Workflows Survey 2017

Thank you for participating in the ECP Workflows Survey for 2017. This survey will help ECP leadership determine directions in workflows development and deployment on exascale systems.

For the purposes of this survey, a workflow is a collection of computational tasks and operations in support of some greater scientific goal. This involves the usage of programs and libraries, including their installation, access, execution, and monitoring. It also involves the management of data, its creation, access, sharing, movement, and life cycles. Workflows involve human and automated interactions both on and off of large-scale computers, as well as among scientific codes and system services, on the login nodes, compute nodes, or combinations. Information about the characteristics of these items will help in the development of workflow systems and the overall exascale ecosystem.

Most questions in the survey have multiply-selectable checkboxes so you can easily provide more than one response, fill in a new answer, or indicate that you don't know. You can also use the follow-up comment fields for further information that you may wish to add, which will help us interpret the survey results. For example, if you have multiple workflows that may be of interest, you may use multiple checkboxes and describe your answers, or take the survey twice, indicating the different workflow of interest in your project description.

Demographics

1. **Which ECP project do you represent?**

   _____

2. **What is your name?**

   _____

3. **What is your email address?**

   _____

4. **What is your role?**
   Check all that apply.

   - [ ] PI
   - [ ] Co-PI
   - [ ] Staff
   - [ ] Postdoc or student
   - [ ] Other: _____

5. **Comment (if any):**

6. **Which current or planned computer systems are of interest to you (by individual system name)?**

Computation

This section considers the distribution of the computational and volatile memory aspects of your workflows.

7. **How many programs or software components may be brought together in one of your workflows?**
   Check all that apply.
   - [ ] 1
   - [ ] 2-4
   - [ ] 5-10
   - [ ] 10+
   - [ ] I don't know
   - [ ] Other:

8. **Comment on program count**

9. **What size jobs make up your workflows?**
   Check all that apply.
   - [ ] Single process, single thread
   - [ ] Single node, multithread
   - [ ] 2-16 nodes
   - [ ] 16-1K nodes
   - [ ] 1K-10K nodes
   - [ ] 10K-100K nodes
   - [ ] 100K+ nodes
   - [ ] I don't know
   - [ ] Other:

10. **Comment on node count**

_____

11. **How long are the tasks in your workflow?**
Check all that apply.

- [ ] 1 second or less
- [ ] 1-99 seconds
- [ ] 100 seconds - 15 minutes
- [ ] 16 minutes - 1 hour
- [ ] 1 hour - 12 hours
- [ ] 12+ hours
- [ ] I don't know
- [ ] Other: _____

12. **Comment on task time**

_____

13. **In a given process, how many threads might you use efficiently?**
Check all that apply.

- [ ] 1
- [ ] 2-7
- [ ] 8-63
- [ ] 64+
- [ ] I don't know
- [ ] Other: _____

14. **Comment on threads**

_____

15. **How much memory can you use effectively on a single node?**
Check all that apply.

- [ ] 1 GB or less
- [ ] 2-7 GB
- [ ] 8-15 GB
- [ ] 16-63 GB
- [ ] 64+ GB
- [ ] I don't know
- [ ] Other: _____

16. **Comment on memory usage**

_____

17. **What accelerators are your code capable of using?**
Check all that apply.

- [ ] None
- [ ] Xeon Phi
- [ ] GPU
- [ ] FPGA
- [ ] I don't know
- [ ] Other: _____

18. **Comment on accelerators**

_____

19. **How many tasks (program or software component invocations) can be part of a single workflow?**
Check all that apply.

- [ ] 1
- [ ] 2-19
- [ ] 20-99
- [ ] 100-999
- [ ] 1000-9999
- [ ] 10,000-999,999
- [ ] 1M - 10M
- [ ] 11M - 100M
- [ ] 101M-1000M
- [ ] 1B+
- [ ] I don't know
- [ ] Other: _____

Data

This section considers persistent data storage and access

20. **How much data is consumed and produced by individual tasks in your workflow?**
Check all that apply.

- [ ] 1MB or less
- [ ] 1-99MB
- [ ] 100-999MB
- [ ] 1GB - 9GB
- [ ] 10 GB - 100 GB
- [ ] 100+ GB
- [ ] I don't know
- [ ] Other: _____

21. **Comment on task data size**

_____

22. **How many files are accessed by the whole workflow?**
Check all that apply.

- [ ] 1-10 in total
- [ ] 11-10,000 in total
- [ ] 1-10 per task
- [ ] 11-100 per task
- [ ] 101+ per task
- [ ] I don't know
- [ ] Other: _____

23. **Comment on file count**

_____

24. **What use cases are you considering for burst buffers?**
Check all that apply.

- [ ] Implicitly - behind the POSIX interfaces
- [ ] Implicitly - behind some I/O library
- [ ] Implicitly - the workflow system should be able to stage my data
- [ ] Explicitly - do my own buffering/staging
- [ ] I don't know what burst buffers are.
- [ ] My workflow cannot benefit from burst buffers.
- [ ] Other: _____

25. **Comment on burst buffers**

_____

26. **What types of code coupling are required by your workflow?**
Check all that apply.

- [ ] Shared-memory, zero-copy
- [ ] Explicit MPI communication
- [ ] Data exchange via other communication library or framework
- [ ] Communication via files
- [ ] Databases
- [ ] None, my tasks are independent.
- [ ] I don't know what code coupling is.
- [ ] Other: _____

27. **Comment on code coupling**

_____

Workflow systems and methods

28. **What workflow systems do you currently use?**

_____

29. **What workflow systems are you considering using?**

_____

30. **What paradigms drive your workflows?**
Check all that apply.

- [ ] Bag of tasks
- [ ] Optimization
- [ ] Uncertainty quantification
- [ ] Machine learning
- [ ] Simulation + analytics
- [ ] Coupled simulations (e.g., multi-physics, multi-scale)
- [ ] Other: _____

31. **Comment on paradigms**

_____

32. **Do your workflows have a human-in-the-loop?**
Check all that apply.

- [ ] Humans have to inspect every task
- [ ] The workflow is fully automated
- [ ] Somewhere in between
- [ ] I don't know
- [ ] Other: _____

33. **Comment on human-in-the-loop**

_____

34. **What external data impacts your workflow?**
Check all that apply.

- [ ] Experiment (e.g., light sources, telescopes)
- [ ] Streams from other computational jobs
- [ ] I don't know
- [ ] Other: _____

35. **Comment on external data**

_____

Workflow programmability

36. **What is your preferred syntax for workflow development?**
Check all that apply.

- [ ] Need a rich programming language
- [ ] Need a standard programming language (shell/Makefile)
- [ ] Embedded in a compiled language (C/C++/Fortran)
- [ ] Embedded in a scripting language (Python, etc.)
- [ ] Static data-flow workflow description (Directed Acyclic Graph of Processes/Tasks)
- [ ] Bag of tasks
- [ ] Visual workflow composition (graphical user interface, GUI)
- [ ] My workflows are written by others
- [ ] I don't know
- [ ] Other: _____

37. **Comment on syntax**

_____

38. **Does your workflow have interoperability requirements?**
Check all that apply.

- [ ] Need to integrate with MPI
- [ ] Need to integrate with the system scheduler
- [ ] Other: _____

39. **Comment on interoperability**

_____

40. **How are your workflows composed?**
Check all that apply.

- [ ] My codes must/may be separate executables
- [ ] My codes must/may be executed as libraries in a framework
- [ ] My codes must/may be run on the same node
- [ ] I don't know
- [ ] Other: _____

41. **Comment on workflow composition**

_____

42. **What are your preferred interfaces for managing simulations and human-in-the-loop interactions?**
Check all that apply.

- [ ] Standard command line interface
- [ ] Notebooks: interactive, sharable scripting
- [ ] IDE or similar environment
- [ ] Other GUI, dashboard, or web interface
- [ ] I don't know
- [ ] Other: _____

43. **Comment on user interfaces**

_____

44. **Do you face challenges in running your workflows on current systems, for example, in launching multiple executables, mapping them to specific nodes, and/or orchestrating communications? If so, please describe.**

_____

_____

_____

_____

_____

Workflow features

45. **How important is the reproducibility of workflows?**
Mark only one oval.

	1	2	3	4	5	6	7	8	9	10	
Not important	◯	◯	◯	◯	◯	◯	◯	◯	◯	◯	Very important

46. **Comment on reproducibility**

_____

47. **How important is robust restart on failure?**
Mark only one oval.

	1	2	3	4	5	6	7	8	9	10	
Not important	◯	◯	◯	◯	◯	◯	◯	◯	◯	◯	Very important

48. **Comment on restart**

_____

49. **Can your entire workflow be checkpointed/restarted?**
Check all that apply.

- [ ] Yes
- [ ] No
- [ ] In part
- [ ] I don't know
- [ ] Other: _____

50. **Comment on checkpoint/restart**

_____

51. **Which of the following describes your provenance needs or capabilities?**
Check all that apply.

- [ ] I collect my own provenance data
- [ ] I need workflow provenance data
- [ ] Log files are fine
- [ ] I would like to use a database system
- [ ] I also need integration with performance analysis tools
- [ ] I don't know what provenance is
- [ ] Other: _____

52. **Comment on provenance**

Free form questions

Other aspects that are less easily categorized

53. **Please describe any requirements not captured above.**

54. **Are there particular tools or technologies you would like to see supported in the ECP stack for your application?**

55. **What are the key challenges (systems, software, administrative) do you face in running your workflows?**

56. **What do you expect from workflows at exascale?**

Data Science & Learning Division
Argonne National Laboratory
9700 South Cass Avenue
Lemont, IL 60439

www.anl.gov

Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC